


DESSINS D'ENFANTS FOR ANALYSTS

VINCENT BEFFARA

ABSTRACT. We present an algorithmic way of exactly computing Belyi functions for hypermaps and triangulations in genus 0 or 1, and the associated dessins, based on a numerical iterative approach initialized from a circle packing combined with subsequent lattice reduction. The main advantage compared to previous methods is that it is applicable to much larger graphs; we use very little algebraic geometry, and aim for this paper to be as self-contained as possible.



Date: April 2, 2015.

This work was started during a stay at the Max Planck Institute and the Hausdorff Center in Bonn, and completed at the Newton Institute in Cambridge; support of these institutions is gratefully acknowledged.

INTRODUCTION

Let T be a triangulation of the sphere \mathbb{S}^2 , *i.e.* a planar graph embedded in the sphere in such a way that all the faces it delimits have degree 3. From T , one can construct a complex structure on the sphere by gluing together equilateral triangles according to the combinatorics of T ; this leads to a manifold with conical singularities at the locations of the vertices of T , and the Riemann uniformization theorem ensures the existence of a conformal bijection Φ_T between this surface and the Riemann sphere \mathbb{C}^* . The image of T by Φ_T is a triangulation of the Riemann sphere with the same combinatorics as T and edges embedded as analytic curves, which is well-defined up to the action of the Möbius group.

While this construction is quite simple, *computing* the embedding explicitly is very difficult in general, as soon as the number of vertices of T is not very small. One of the goals of this paper is to describe an algorithm to do it automatically, in a “semi-numerical” way (that is, going through an approximate, numerical stage to obtain an exact, explicit outcome).

Assume in addition that all the vertices of T have even degree. This is equivalent to two other properties of the triangulation: first, its faces can be partitioned into two sets in such a way that a face in one set is only adjacent to faces in the other set; this is usually done graphically by coloring one half of the faces white and the other half black. Second, the vertices of T can be partitioned into three classes in such a way that along the boundary of each face, exactly one vertex of each class appears.

Choosing the triangles used to build our Riemann surface to be the two hemispheres of the Riemann sphere separated by the extended real line, with vertices at 0, 1 and ∞ , from the data of T we obtain a covering $\pi_T : \mathbb{C}^* \rightarrow \mathbb{C}^*$ that is ramified only above $\{0, 1, \infty\}$ (such a map is known as a *Belyi function*), and the embedding of T in \mathbb{C}^* can be seen as the preimage of $\mathbb{R} \cup \{\infty\}$ by π_T , the three classes of vertices being the preimages of 0, 1 and ∞ respectively. So, the question of computing the embedding turns into that of computing π_T .

The covering π_T is meromorphic, and in the case of the sphere it means that it has to be a rational function; computing it then means computing the location of its zeros and its poles, or equivalently the coefficients of its numerator and denominator. It is always possible to choose the embedding (*via* a Möbius transformation) in such a way that all these numbers are algebraic, and what we are looking for is their specification as roots of explicit polynomials with integer coefficients. This has been done in many cases before, cf. for example [5, 6, 13, 15], tracing back to the original work of Klein on the icosahedron and its link with fifth-degree equations [10]; a recent work with a similar goal as ours, but very different methods, is [3].

The same construction can be performed from a triangulation of any orientable surface, and leads to a Riemann surface M_T of the same genus together with a ramified covering $\pi_T : M_T \rightarrow \mathbb{C}^*$; this surface in turn can be seen as a (smooth irreducible projective) algebraic curve, and Belyi’s theorem (see below for a precise statement) ensures that this curve is defined over $\bar{\mathbb{Q}}$. A similar goal as before is then to identify this algebraic curve explicitly, by specifying the coefficients of its equation as roots of explicit integer polynomials, as well as those of the covering.

In this paper, besides the sphere we will only discuss the simpler case of genus 1, where the Riemann surface built from T can always be uniformized into a torus $\mathbb{T} = \mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$ for some modulus τ (depending on T) in the complex upper-half plane. The previous problem becomes the exact computation of τ and of the covering $\pi_T : \mathbb{T} \rightarrow \mathbb{C}^*$. Lifting π_T to the universal

cover of \mathbb{T} , this means that we are looking for an elliptic function rather than for a rational one as above, but apart from that the situation remains very similar in principle, in that we are looking for a finite collection of algebraic numbers. The higher-genus situation should be amenable to a similar treatment as the one we describe here, but the implementation would have to be significantly more complex.

The usual ways that the question has been addressed previously are mostly algebraic in nature: one can write the conditions that a function has to satisfy in order to be a Belyi function as a system of polynomial equations, and try to solve it exactly from the start. This works well for small examples, but the combinatorial complexity grows very quickly and even moderately large cases, especially in positive genus, are beyond the reach of computer algebra systems. A recent proposal by Bartholdi et. al. [3] for instance quotes a computing time of 15 minutes for a triangulation of the sphere with 15 vertices.

Here instead we base our approach on a numerical point of view, and the central step is an implementation of the Newton algorithm to get an arbitrary precision approximation of the algebraic numbers we are looking for, initialized with a configuration obtained using circle packings; this can then be combined with a lattice-reduction algorithm to identify the corresponding integer polynomials. For comparison, for the same example as used in [3], getting 100 digits of precision for all parameters (which is more than enough to obtain an exact solution) takes less than a second.

Our initial motivation for this work was related to probabilistic conjectures on the conformal structure of large random planar maps; testing these conjectures numerically without computing the embedding exactly is difficult, because in this setup approximations come from many sources: from the randomness of the map, from the fact that it is of finite size, and also from the embedding approximation itself, so determining the embedding exactly is of practical importance. Computing the embedding algebraically being beyond reach of the previous methods for the numbers of vertices that are relevant in this setup, we were led to look for a more numerical approach which is the focus of this paper; we will apply it to random triangulations in a subsequent work.

The remainder of this paper is organized as follows: section 1 contains the formal definitions of the objects that we are using, section 2 gives a statement of Belyi's theorem and is there mostly to keep the argument self-contained, and sections 3 and 4 describe our approach in detail and provide proofs of convergence in the cases of genus 0 and 1, respectively. Section 5 then gives a few examples of application.

The source code implementing the algorithms, and used to produce all the examples and pictures in this paper, is publicly available at the following address:

<http://github.com/vbeffara/Simulations>

1. DEFINITIONS

1.1. Maps and hypermaps. Our starting point will always be a graph drawn on a given topological surface, but seen as a combinatorial object, *i.e.* up to homeomorphisms of the surface. All the following definitions are classical, but we still give them for sake of self-containedness. Much more can be found for instance in the book of Lando and Zvonkin [11].

Definition. A (finite) *graph* is a pair $G = (V, E)$ where V is seen as the set of *vertices* of G and $E \subset V^2$ as the set of its *edges*. We will always consider *undirected* graphs, for which E is symmetric: $(x, y) \in E \iff (y, x) \in E$, and assume the absence of *loops*, i.e. of edges of the form (x, x) linking a vertex to itself.

Definition. An *embedding* of an undirected graph $G = (V, E)$ in an oriented surface M is the data of a collection $(x_v)_{v \in V}$ of pairwise distinct points of M , and a collection $(\gamma_{x,y})_{(x,y) \in E}$ of continuous simple curves on M , with $\gamma_{x,y} : [0, 1] \rightarrow M$ satisfying $\gamma(0) = x$ and $\gamma(1) = y$ and the symmetry condition $\gamma_{x,y}(t) = \gamma_{y,x}(1 - t)$, and such that curves corresponding to distinct edges have disjoint images except possibly for their endpoints. The embedding is *proper* if every connected component of the complement of the union of all the images of the $\gamma_{x,y}$ (in other words, each *face*) is homeomorphic to a disk.

Definition. A *map* is an equivalence class of proper embeddings of a graph G in a surface M , where two embeddings are identified if they are conjugated by a homeomorphism of M . When M is the Riemann sphere \mathbb{C}^* , we will speak about a *planar map*; when M is a two-dimensional torus, about a *toroidal map*. If each of the faces has exactly 3 edges along its boundary, the map is called a *triangulation* of M .




FIGURE 1. A planar map with 5 vertices, 6 edges and 3 faces (including the outer one in this representation).

A map is in fact a combinatorial structure, in the sense that there are only finitely many maps consisting in embeddings of a given finite graph in a given surface; it can be specified by ordering the neighbors of each of the vertices of the graph in a cyclic order. This remark allows for the definition of the *dual map* of a map: its (dual) vertices are in bijection with the faces of the map; two dual vertices will be declared adjacent if the corresponding faces of the initial map share an edge, and the neighbor ordering in the dual map is then given by the cyclic order around the corresponding face of the primal map.

A very related structure, which makes computer implementations easier to manage in practice and is of algebraic relevance, is the following:

Definition. Given a positive integer n , a *hypermap* of size n is a triple $(\sigma, \alpha, \varphi)$ of permutations of a set of n elements satisfying the compatibility relation $\sigma\alpha\varphi = \text{id}$.

This can be seen as a generalization of the notion of map: indeed, restricting to the case where the cycles of α are all transpositions, one can interpret the domain of the permutations as the set of all the half-edges of the map, the cycles of σ as its vertices (reading the half-edges

incident to that vertex in counterclockwise order), those of α as its edges (which are pairs of half-edges) and those of φ as its faces (reading the half-edges emanating from the vertices along that face in direct order) — see Figure 2 for an example, which is much clearer than any formal description would be.




FIGURE 2. Numbering of the half-edges of the map from Figure 1 to obtain its hypermap representation: in this case, we obtain $\sigma = (0\ 2)(1\ 6\ 4)(3\ 5\ 10)(7\ 8)(9\ 11)$, $\alpha = (0\ 1)(2\ 3)(4\ 5)(6\ 7)(8\ 9)(10\ 11)$ and $\varphi = (0\ 4\ 3)(1\ 2\ 10\ 9\ 7)(5\ 6\ 8\ 11)$. Following the usual convention, the label of each half-edge is to its left.

The maps given as examples in this paper (for instance in Appendix 5) are described in this form, with the n -element set chosen as $\{0, \dots, n-1\}$. The triple corresponding to the dual map is then $(\varphi^{-1}, \alpha^{-1}, \sigma^{-1})$ — and obviously here $\alpha^{-1} = \alpha$, but this way of denoting it makes the notation clearer. Conversely, hypermaps can also be seen as a particular case of maps:

Definition. A graph is called *bipartite* if its vertex set can be partitioned into two disjoint subsets, in such a way that each vertex in one of them is only adjacent to vertices in the other one. A map is bipartite if the associated graph is bipartite; if the map is proper, this is equivalent to saying that its faces all have an even number of edges along their boundaries.

It is customary to refer to the partition of the vertices of a bipartite map into *black* and *white* vertices. Such a map can then be encoded as a hypermap, where the cycles of σ (resp. α , φ) correspond to the black vertices (resp. white vertices, faces). In this case, the cycles of φ have a length equal to a half of the number of edges on the boundary of the corresponding faces of the map. If every white vertex has degree 2, then the black vertices themselves form a graph with the same hypermap representation.

One can compose these two constructions, starting from a map, seeing it as a hypermap where the cycles of α have length 2, and then seeing this hypermap as a bipartite map where the black vertices correspond to the vertices of the initial map and the white ones (with degree 2) to its edges. More graphically, this is equivalent to adding one vertex on each edge of the initial map, as shown in Figure 3.

The last identification that we will need is with particular triangulations:

Definition. A triangulation of a surface is called *tripartite* if its vertex set can be partitioned into 3 disjoint subsets, in such a way that each of its faces has one vertex of each of these subsets along its boundary. This is equivalent to saying that its dual map is bipartite, and that all its vertices have even degree.




FIGURE 3. The bipartite map corresponding to the hypermap in Figure 2.

The three vertex sets will be represented here as black, white and red; in [11], vertices of the third kind are represented as asterisks. To each tripartite triangulation correspond 3 bipartite maps, obtained by keeping the vertices in two out of the three subsets in the definition and the edges connecting them. Conversely, every proper bipartite map can be obtained this way, and another way of stating this is that every proper bipartite map can be completed into a tripartite triangulation of the same surface by adding a vertex inside each face and connecting it to the original vertices on the boundary of that face.

Combining the previous remarks, every (proper) planar map can be refined into a tripartite triangulation by adding one vertex on each of its edges and one vertex inside each of its faces, connecting them in the natural way. If $(\sigma, \alpha, \varphi)$ is the hypermap corresponding to the initial map, then the cycles of each of these permutations are in bijection with the vertices in one of the three subsets in the partition of the tripartite triangulation. We will refer to this triangulation as the *tripartite refinement* of the map — see Figure 4. In the diagrams of Section 5, the edges incident to red vertices are not represented in order to make the combinatorics more readable.




FIGURE 4. The tripartite refinement of the planar map from Figure 1. One can check that the degree of each red vertex is equal to twice the number of edges of the corresponding face in the initial map.

1.2. Constellations. Our main objects of study in this paper are holomorphic maps defined between Riemann surfaces, but it is convenient as a way of encoding Belyi maps to introduce a more discrete kind of object. In this whole section, M will denote a fixed Riemann surface. Our first definition is that of a decorated point of M :

Definition. A *star* on M is a pair $\mathbf{z} = (z, d)$ where $z \in M$ (the *location*) and d is a positive integer (the *multiplicity*).

The idea is that a star will specify the behavior of a map at its location (typically d would be the degree of ramification at z). A polynomial can be given as a finite sequence of stars located at its roots, and a rational fraction as two such sequences (for the roots and the poles). To match later with Belyi's theorem, we will need just a little more information:

Definition. A *constellation* on M is a triple $\mathcal{C} = (Z, P, O)$ where $Z = \{\mathbf{z}_i\}_{1 \leq i \leq n_z}$, $P = \{\mathbf{p}_i\}_{1 \leq i \leq n_p}$ and $O = \{\mathbf{o}_i\}_{1 \leq i \leq n_o}$ are three finite (potentially empty) sets of stars, which we will refer to respectively as the *zeros*, *poles* and *ones* of \mathcal{C} . A constellation is *non-degenerate* if the locations of all involved stars are pairwise distinct (which we will always implicitly assume except otherwise mentioned).

If $f : M \rightarrow \mathbb{C}^*$ is a holomorphic covering of the Riemann sphere, it comes naturally with a constellation $\mathcal{C}_f = (Z_f, P_f, O_f)$ where Z_f lists the preimages of 0, P_f those of ∞ and O_f those of 1, and the multiplicities are the corresponding orders of ramification. Under the assumptions that f only ramifies over $\{0, 1, \infty\}$, f is known as a *Belyi function*, and \mathcal{C}_f characterizes f uniquely; this relation can be made much more explicit in a few cases, which we describe now.

1.2.1. Polynomials on \mathbb{C} . If $f \in \mathbb{C}[X]$ is a polynomial with complex coefficients and degree d , it is specified uniquely by the collection of its roots and their multiplicities, and its leading coefficient: if its constellation is given as

$$\mathcal{C}_f = (\{(z_i, d_i)\}, \{(\infty, d)\}, \{(o_i, d'_i)\})$$

(with $\sum d_i = d$) then f can be recovered as

$$f(z) = \lambda \prod (z - z_i)^{d_i}$$

where λ is chosen to make the value of f equal to 1 at the points of O_f . In this setup, f is a Belyi function if and only if $\sum d'_i = d$, in which case

$$f(z) - 1 = \lambda \prod (z - o_i)^{d'_i} \quad \text{and} \quad f'(z) = \lambda d \prod (z - z_i)^{d_i-1} \prod (z - o_i)^{d'_i-1}.$$

The polynomial f is then known as a *Shabat polynomial*.

1.2.2. Rational fractions on \mathbb{C} . Let f now be a (ramified) covering of the Riemann sphere by itself: it has to be a rational fraction with complex coefficients, and can be written as $f(z) = p(z)/q(z)$ where p (resp. q) is a polynomial of degree d_1 (resp. d_2). Again $Z_f = \{(z_i, d_i)\}$ and $P_f = \{(p_i, d'_i)\}$ can be written explicitly in terms of the roots of p and q respectively, and ∞ will occur in Z_f , in P_f or in neither according to the relative values of d_1 and d_2 . f can then be similarly reconstructed as

$$(1) \quad f(z) = \lambda \frac{\prod (z - z_i)^{d_i}}{\prod (z - p_i)^{d'_i}}$$

(where the products omit the potential term at ∞ , and where as before λ is chosen so that f takes value 1 at the o_i). Once more, the constellation \mathcal{C}_f corresponds to a Belyi function if and only if the sums of multiplicities are the same for Z_f , P_f and O_f .

1.2.3. Functions on a complex torus. The case of genus 1 can also be made quite explicit. We will assume in this whole section that M is the complex torus with periods 1 and τ where $\Im(\tau) > 0$, in other words

$$M = \mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z}).$$

We will abuse notation by writing the locations of stars as complex numbers. Here, M comes with a natural covering $\pi : \mathbb{C} \rightarrow M$ and if $f : M \rightarrow \mathbb{C}^*$ is holomorphic, it can be lifted as a periodic meromorphic function $\hat{f} : \mathbb{C} \rightarrow \mathbb{C}^*$ (an *elliptic function*).

The situation is a little more rigid than before, in the sense that some relations are automatically satisfied by the zeros and poles of f . More specifically, if $Z_f = \{(z_i, d_i)\}$ and $P_f = \{(p_i, d'_i)\}$ then

$$(2) \quad \sum d_i = \sum d'_i \quad \text{and} \quad \sum d_i z_i \equiv \sum d'_i p_i$$

(the second relation being meant mod $\mathbb{Z} + \tau\mathbb{Z}$). The reconstruction of f from these can be made in a similar fashion as in the case of rational fractions: if the p_i and z_i are chosen in such a way that $\sum d_i z_i = \sum d'_i p_i$ (this time as complex numbers in \mathbb{C}), then f can be obtained as

$$(3) \quad f(z) = \lambda \frac{\prod \zeta(z - z_i)^{d_i}}{\prod \zeta(z - p_i)^{d'_i}}$$

in terms of the Weierstrass function ζ with (quasi-)periods 1 and τ . It is also the case that f can be written as a polynomial in the Weierstrass function \wp with the same periods and its derivative \wp' , or equivalently as a rational function on the associated elliptic curve, but although that is the “right” way to reconstruct f in terms of algebraic geometry, the coefficients of this polynomial depend on the constellation in a way that is less transparent and does not exhibit the same similarity with the previous two cases.

1.2.4. A few additional remarks. In the case of Belyi functions, it is always the case that the sums of multiplicities is the same for Z_f , P_f and O_f , being equal to the degree of the covering; and the genus of the surface M is related to these multiplicities by the Riemann-Hurwitz formula. It is therefore natural to define the following:

Definition. A constellation $\mathcal{C} = (Z, P, O)$ is *balanced* if it satisfies

$$\sum_{(w,d) \in Z} d = \sum_{(w,d) \in P} d = \sum_{(w,d) \in O} d =: N(\mathcal{C}).$$

If \mathcal{C} is balanced, its *genus* is given by

$$g(\mathcal{C}) := 1 + \frac{N(\mathcal{C}) - \#(Z \cup P \cup O)}{2}.$$

With these notations, if $f : M \rightarrow \mathbb{C}^*$ is a Belyi function, then its constellation \mathcal{C}_f is balanced and its genus is that of M . As we saw previously, if this genus is 0 or 1, it is then possible to explicitly reconstruct f from \mathcal{C} . Besides, if $(w, d) \in O_f$, then d gives the degree of ramification of f at w . This can be used in the other direction to characterize the constellations corresponding to a Belyi function:

Definition. Let $\mathcal{C} = (Z, P, O)$ be a balanced constellation of genus 0 or 1: we will say that it is *exact* if there exists $\lambda \in \mathbb{C}$ satisfying the following conditions. Let f be defined according to either (1) or (3) (depending on the genus — and assume that (2) is satisfied in the case of genus 1). Then for every $(w, d) \in O$:

- $f(w) = 1$;
- for every $k \in \{1, \dots, d-1\}$, $f^{(k)}(w) = 0$.

Such a value of λ is obviously unique, and will be called the *canonical normalization* of \mathcal{C} .

1.3. Drawing a constellation. Let M be of genus 0 or 1 and $\mathcal{C} = (Z, P, O)$ be an exact constellation on M ; let $f : M \rightarrow \mathbb{C}^*$ be the corresponding covering. This setup allows to draw two natural, related structures on M .




FIGURE 5. Uniformization of a periodic triangulation of 90 vertices. The color corresponds to the sign of the imaginary part of the covering map.

1.3.1. As a triangulation. First, one can look at the tripartite map on M with vertices at $Z \cup P \cup O$, and edges given as the preimages of the three intervals $(-\infty, 0)$, $(0, 1)$ and $(1, +\infty)$ of \mathbb{R} by f . This is a triangulation of M , and its dual is a bipartite map with all vertices of degree 3. Most of the pictures in this article are drawn in this setup (see *e.g.* the first page

for an example in genus 0, or Figure 5 in genus 1), and the faces are colored according to the sign of the imaginary part of f on them.

1.3.2. *As a hypermap.* Another way to draw \mathcal{C} is as a hypermap, where Z consists in the black vertices, P in the white vertices, and the edges are the preimages of the segment $[0, 1]$ by f . The combinatorics of the hypermap (the triple of permutations (σ, α, ϕ) in the language of the previous section) is equivalent to that of the triangulation above, and is completely determined by the constellation. It is this particular way of drawing this hypermap on M which is usually referred to as a *dessin d'enfant*. Even though the information is theoretically the same as before, it is computationally more difficult to obtain visually satisfactory pictures this way.

2. BELYI'S THEOREM

We give in this section an extremely limited introduction to Belyi's theorem, for the benefit of readers who are not familiar with the topic. For much more, we refer to [9] and references therein.

2.1. **Preliminary remarks.** For fixed balanced multiplicities, the space of constellations is of complex dimension $\#(Z \cup P \cup O)$. The number of equations to be satisfied for a constellation \mathcal{C} to be exact is counted a bit differently depending on the genus:

- If $g = 0$, it is the number of conditions coming from the definition, which is equal to $N(\mathcal{C}) - 1$;
- If $g = 1$, the relation (2) needs to be counted as well and the total number is $N(\mathcal{C})$.

The difference δ between the number of variables and the number of conditions is then equal to 3 in genus 0, and to 0 in genus 1; in particular it depends only on the genus.

Before we can make use of this computation, one remark is in order. In the general setting, automorphisms of M act in a natural way on constellations on M by mapping locations while preserving multiplicities. This preserves the property of being exact; it is natural to identify two constellations conjugated by such an automorphism, or at least to classify them up to automorphism.

In genus 0, the group of automorphisms of the Riemann sphere has complex dimension 3, which is equal to δ . This means that we should expect the number of exact constellations of given multiplicities, up to automorphisms of M , to be finite.

In genus 1, the group of automorphisms of M has complex dimension 1 (in the generic case it consists purely of translations), which is one more than δ . This means that in the general case we should expect the existence of no exact constellation with given multiplicities, but that for finitely many values of τ there should exist finitely many exact constellations.

Of course, all the preceding remarks are at the heuristic level and making them formal would imply controlling degeneracies and genericity, which would likely be quite difficult to do at this point. Nevertheless, as what follows will show, they do give the right predictions.

2.2. **Analytic statement.** Let T be a map with triangular faces, of genus 0 or 1. It can be made into a Riemann surface M_T by gluing together equilateral triangles according to its combinatorics; and this surface can be uniformized to either the Riemann sphere, or to a complex torus with periods 1 and τ_T , where τ_T is uniquely determined (up to $SL_2(\mathbb{Z})$ action) by T . It is a natural question, and the origin of the work presented here, to in the latter

case determine the value of τ_T from the discrete data of the map T , either exactly or at least numerically (which, if done with enough precision, suffices to obtain an exact value).

If the vertices of T all have even degree (which is in particular the case if T is associated to a hypermap), they can be seen as stars with multiplicity half of their degree and it is always possible to split them into three disjoint sets (or types) Z , P and O in such a way that each of the faces contains one vertex in each of these sets. One can then use the uniformization obtained above to map M_T onto the Riemann sphere, each triangle being mapped to either the upper or lower hemisphere according to the order in which the vertex types occur on its boundary. This construction leads to a covering $f_T : M_T \rightarrow \mathbb{C}^*$ which is a Belyi function, and (Z, P, O) is then the exact constellation associated to f_T ; see Figure 5 for an illustration in the case of genus 1.

This gives a justification to the predictions at the end of the previous section: since each exact constellation gives rise to a triangulation, it can be obtained (up to automorphism) from that same triangulation. In turn, since the number of triangulations of a given size is finite, this implies that the number of exact constellations of given multiplicities, counted up to automorphism, is itself finite, and so is the set of values of τ such that $\mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$ supports an exact constellation of a given size.

This approach is very satisfactory at the theoretical level; unfortunately, it doesn't lend itself very well to explicit computations, and performing this program numerically in a direct way would be difficult.

2.3. Number-theoretic statement. There is a deep connection between the complex structure introduced above and the arithmetic properties of the underlying Riemann surface, having its origin in the following equivalence:

Theorem 1 (Belyi). *Let C be a non-singular algebraic curve: then there exists a ramified covering $C \rightarrow \mathbb{C}^*$, ramified only above $\{0, 1, \infty\}$, if and only if C is defined over $\bar{\mathbb{Q}}$.*

We will mostly be concerned with the “only if” part in what follows; notice that when we build a Riemann surface from gluing equilateral triangles, it automatically comes with a constellation (Z, P, O) and such a covering, namely the function f constructed above, sending the points in Z (resp. P , O) to 0 (resp. ∞ , 1) and unramified above the other points.

What it means for us is that the torus $\mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$ built from a triangulation of genus 1, seen as an elliptic curve, is defined over $\bar{\mathbb{Q}}$. In particular, the modulus τ can theoretically be computed explicitly. In practice, this seems to be possible only in very few cases, and in each of them, requires the use of particular symmetries of the triangulation.

3. NUMERICAL COMPUTATION OF BELYI MAPS: GENUS 0

We now turn to the main point of this paper, namely a semi-numerical strategy to exactly compute the branched covering associated to a given hypermap. As argued above, this can be reduced to determining the exact constellation associated to the map, which lends itself to numerical approaches. We start with the case of hypermaps drawn on the sphere, which benefits from lighter notation while retaining most of the features of our approach.

Fix three tuples of positive integers, $d^0 = (d_i^0)_{1 \leq i \leq n_0}$, $d^\infty = (d_i^\infty)_{1 \leq i \leq n_\infty}$ and $d^1 = (d_i^1)_{1 \leq i \leq n_1}$ with the same sum N , satisfying the genus-0 condition

$$n_0 + n_\infty + n_1 = N + 2,$$

and let $\mathfrak{C} = \mathfrak{C}_{d^0, d^\infty, d^1}$ be the space of all non-degenerate constellations on \mathbb{C}^* with signature (d^0, d^∞, d^1) (which are all balanced and of genus 0). \mathfrak{C} is a complex manifold of dimension $N + 2$. If $\mathcal{C} = (Z, P, O) \in \mathfrak{C}$, let

$$f_{\lambda, \mathcal{C}} : z \mapsto \lambda \prod_{(w, d) \in Z} (z - w)^d \prod_{(w, d) \in P} (z - w)^{-d}$$

be defined as above (omitting the star at ∞ in the products, if any), and consider the map $\Phi : \mathbb{C} \times \mathfrak{C} \rightarrow \mathbb{C}^N$ defined by

$$\Phi(\lambda, \mathcal{C}) := (f(o_1) - 1, f'(o_1), \dots, f^{(d_1^1-1)}(o_1), f(o_2) - 1, f'(o_2), \dots, f^{(d_{n_1}^1-1)}(o_{n_1})),$$

where to lighten notation we let $f = f_{\lambda, \mathcal{C}}$ and $\mathcal{C} = (Z, P, O)$ with $O = ((o_i, d_i^1))$. The constellation \mathcal{C} is exact if and only if there exists $\lambda \in \mathbb{C}$ such that $\Phi(\lambda, \mathcal{C}) = (0, \dots, 0)$.

To normalize the map embeddings, let \mathfrak{C}^* be the set of all constellations $\mathcal{C} = (Z, P, O) \in \mathfrak{C}$ such that $z_1 = 0$, $p_1 = \infty$ and $o_1 = 1$. \mathfrak{C}^* is a manifold of complex dimension $N - 1$ which can be seen as an open subset of \mathbb{C}^{N-1} by listing the locations of all the stars except for z_1 , p_1 and o_1 . The restriction $\Phi^* : \mathbb{C} \times \mathfrak{C}^* \rightarrow \mathbb{C}^N$ can thus be seen as a map from an open subset of \mathbb{C}^N to \mathbb{C}^N , which is clearly analytic in all its variables.

From the previous discussion, the set of preimages of $(0, \dots, 0)$ by Φ^* is finite, and each of its elements corresponds to an exact constellation and to the associated hypermap and covering. From this, the general structure of the algorithm is rather clear: starting from a tripartite hypermap M and the associated triangulation of genus 0, to find its constellation \mathcal{C}_M we will perform the following.

- Step 1: find an approximation \mathcal{C}_0 of \mathcal{C}_M ;
- Step 2: refine the approximation via an iterative scheme, using \mathcal{C}_0 as a starting point;
- Step 3: verify that the solution obtained is the right one (topologically); otherwise go back to step 1 with better precision;
- Step 4: identify the coefficients as algebraic numbers;
- Step 5: verify that the solution obtained is the right one (algebraically).

We now describe each of these steps in some detail, focusing more on the first two which constitute the main contribution of the section.

3.1. Step 1: Approaching the complex structure. To initialize the algorithm, we need to get an approximation of \mathcal{C}_M . This can be done by approximating the complex structure derived from M , and there are several options available here. From the complex analytical point of view, one natural possibility would be to start from an arbitrary embedding of M into the sphere with straight edges, and to identify the uniformizing map as a solution to a Beltrami equation with a Beltrami differential taken to be constant on each of the faces. This has been implemented in a few cases but is quite involved [8], so we chose a simpler way, following [7].

Given a triangulation T of the sphere, there exists a *circle packing* with the combinatorics of T , namely a collection (C_v) of circles indexed by the vertex set of T , such that the associated disks have disjoint interiors and such that C_v and C_w are tangent if and only if v and w are adjacent vertices of T . The circle packing is unique up to Möbius transformations, and it provides an embedding of T in which the vertices are mapped to the centers of the circles and the edges are each the union of two circle radii; in particular, if T is tripartite it provides also a balanced constellation $\mathcal{C}_{cp}(T)$. For general reference on circle packings, we refer the reader to [14]

Our first approximation of \mathcal{C}_M will simply be $\mathcal{C}_{cp}(M)$. In practice, it turns out to often be good enough for our purposes, but it is not always the case. To get a better one, we follow the strategy of [7] and refine M in the following way. Given a triangulation, one can replace each of its faces by 4 new triangles by inserting a vertex at the mid-point of each of its edges — see Figure 6. Iterating the procedure starting from $M_0 := M$, one gets a sequence $(M_n)_{n \geq 0}$ of triangulation, which are all tripartite. The main result is then the following:




FIGURE 6. Successive refinements of a triangulation.

Theorem 2 (Theorem 4.8 in [7]). *Under appropriate normalization, $\mathcal{C}_{cp}(M_n)$ converges to \mathcal{C}_M as $n \rightarrow \infty$.*

This is exactly what we needed: we now have a sequence of constellations which approaches \mathcal{C}_M and will serve as starting points for the iterative method in the next step. We first run steps 2 and 3 starting from $\mathcal{C}_{cp}(M_0) = \mathcal{C}_{cp}(M)$, then if that fails from $\mathcal{C}_{cp}(M_1)$, $\mathcal{C}_{cp}(M_2)$, \dots until the first one that succeeds — which is guaranteed to happen after finitely many subdivisions by Theorems 2 and 3.

3.2. Step 2: Newton's method. This part is quite standard, but as far as we can tell it has not been implemented in this setup before: one can apply a multi-dimensional Newton's method to build a sequence of constellations which converges to \mathcal{C}_M , provided one starts close enough to it.

More specifically, still identifying \mathfrak{C}^* with an open subset of \mathbb{C}^{N-1} , let $\Psi : \mathbb{C} \times \mathfrak{C}^* \rightarrow \mathbb{C} \times \mathfrak{C}^*$ be defined by

$$\Psi((\lambda, \mathcal{C})) := (\lambda, \mathcal{C}) - J_{\Phi}((\lambda, \mathcal{C}))^{-1} \Phi((\lambda, \mathcal{C}))$$

whenever the Jacobian J_{Φ} is nonsingular at (λ, \mathcal{C}) . The main statement of this section is the convergence of Newton's method:

Theorem 3. *Let λ_M be the canonical normalization of \mathcal{C}_M : there exists a neighborhood U of $(\lambda_M, \mathcal{C}_M)$ in $\mathbb{C} \times \mathfrak{C}^*$ satisfying $\Psi(U) \subset U$, on which J_{Φ} is everywhere nonsingular, and such that, whenever $(\lambda, \mathcal{C}) \in U$, the sequence defined inductively by $x_0 = (\lambda, \mathcal{C})$ and $x_{n+1} = \Psi(x_n)$ converges to $(\lambda_M, \mathcal{C}_M)$ as $n \rightarrow \infty$.*

It will be more convenient to set up the computations in the proof in terms of logarithmic derivatives, so let $g_{\lambda, \mathcal{C}}(z)$ be a determination of $\log(f_{(\lambda, \mathcal{C})}(z))$ that is jointly continuous in (λ, \mathcal{C}) in a neighborhood V of $(\lambda_M, \mathcal{C}_M)$ and in z in neighborhoods of each of the o_i , and chosen such that $g_{(\lambda_M, \mathcal{C}_M)}(o_i) = 0$ for every $i \in \{1, \dots, n_1\}$. Let Λ be the counterpart of Φ defined in V in terms of g instead of f , namely:

$$\Lambda(\lambda, \mathcal{C}) := (g(o_1), g'(o_1), \dots, g^{(d_1^1-1)}(o_1), g(o_2), g'(o_2), \dots, g^{(d_{n_1}^1-1)}(o_{n_1})).$$

Lemma 4. *There is a constant $C \in \mathbb{C} \setminus \{0\}$, depending only on the signature (d^0, d^∞, d^1) , such that the Jacobian determinant of Λ at $(\lambda_M, \mathcal{C}_M)$ is given by*

$$\frac{C \prod_{1 \leq i < j \leq n_1} (o_j - o_i)^{\tilde{d}_i^1 \tilde{d}_j^1} \cdot \prod_{2 \leq i < j \leq n_0} (z_j - z_i) \cdot \prod_{2 \leq i < j \leq n_\infty} (f_j - f_i) \cdot \prod_{i=2}^{n_0} \prod_{j=2}^{n_\infty} (f_j - z_i)}{\lambda \prod_{i=1}^{n_1} \left(\prod_{j=2}^{n_0} (z_j - o_i)^{\tilde{d}_i^1} \cdot \prod_{j=2}^{n_\infty} (f_j - o_i)^{\tilde{d}_i^1} \right)},$$

where $\tilde{d}_i^1 := d_i^1 - 1 + \delta_i^1$. In particular, $\det J_\Lambda((\lambda_M, \mathcal{C}_M)) \neq 0$.

Proof. We first need to compute the entries of the Jacobian matrix $J_\Lambda((\lambda_M, \mathcal{C}_M))$. The iterated derivatives of g are given for $k > 0$ by

$$\frac{(-1)^{k+1}}{(k-1)!} g^{(k)}(w) = \sum_{i=1}^{n_0} \frac{d_i^0}{(w - z_i)^k} - \sum_{i=1}^{n_\infty} \frac{d_i^\infty}{(w - z_i)^k}$$

so the partial derivatives are all explicit:

$$\begin{aligned} \partial_\lambda g^{(k)}(o_i) &= \frac{1}{\lambda} \delta_k^0, & \partial_{o_j} g^{(k)}(o_i) &= g^{(k+1)}(o_i) \delta_i^j, \\ \partial_{z_j} g^{(k)}(o_i) &= \frac{d_j^0 (-1)^{k+1} k!}{(o_i - z_j)^{k+1}}, & \partial_{f_j} g^{(k)}(o_i) &= \frac{d_j^\infty (-1)^{k+1} k!}{(o_i - f_j)^{k+1}}. \end{aligned}$$

In addition, at the point $(\lambda_M, \mathcal{C}_M)$ we know that $g^{(k)}(o_i) = 0$ as soon as $k < d_i^1$, so up to terms depending only on the signature (and which are products of $(-1)^k$, d_j -s and factorials), the Jacobian determinant is proportional to

$$\begin{vmatrix} 1/\lambda & 0 \dots & 0 & 1/\lambda & 0 \dots & 0 & \dots & 1/\lambda & 0 \dots & 0 \\ (o_1 - z_2)^{-1} & \dots & (o_1 - z_2)^{-d_1^1} & (o_2 - z_2)^{-1} & \dots & (o_2 - z_2)^{-d_2^1} & \dots & (o_{n_1} - z_2)^{-1} & \dots & (o_{n_1} - z_2)^{-d_{n_1}^1} \\ \vdots & & \vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots \\ (o_1 - z_{n_0})^{-1} & \dots & (o_1 - z_{n_0})^{-d_1^1} & (o_2 - z_{n_0})^{-1} & \dots & (o_2 - z_{n_0})^{-d_2^1} & \dots & (o_{n_1} - z_{n_0})^{-1} & \dots & (o_{n_1} - z_{n_0})^{-d_{n_1}^1} \\ (o_1 - f_2)^{-1} & \dots & (o_1 - f_2)^{-d_1^1} & (o_2 - f_2)^{-1} & \dots & (o_2 - f_2)^{-d_2^1} & \dots & (o_{n_1} - f_2)^{-1} & \dots & (o_{n_1} - f_2)^{-d_{n_1}^1} \\ \vdots & & \vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots \\ (o_1 - f_{n_\infty})^{-1} & \dots & (o_1 - f_{n_\infty})^{-d_1^1} & (o_2 - f_{n_\infty})^{-1} & \dots & (o_2 - f_{n_\infty})^{-d_2^1} & \dots & (o_{n_1} - f_{n_\infty})^{-1} & \dots & (o_{n_1} - f_{n_\infty})^{-d_{n_1}^1} \\ 0 & \dots & 0 & 0 & \dots & 1 & \dots & 0 & \dots & 0 \\ \vdots & & \vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots \\ 0 & \dots & 0 & 0 & \dots & 0 & \dots & 0 & \dots & 1 \end{vmatrix}$$

where the columns are the components of Λ and the lines correspond to partial derivatives with respect, in order, to λ , the $(z_i)_{i \geq 2}$, the $(f_i)_{i \geq 2}$ and the $(o_i)_{i \geq 2}$ (remember that we normalized all constellations to have $z_1 = 0$, $o_1 = 1$ and $f_1 = \infty$, so they do not appear as variables here, and that this makes the matrix square). Developing the determinant along its last $(n_1 - 1)$

lines shows that it is equal to

$$\begin{vmatrix} 1/\lambda & 0 \cdots & 0 & 1/\lambda & 0 \cdots & 0 & \cdots & 1/\lambda & 0 \cdots & 0 \\ (o_1 - z_2)^{-1} & \cdots & (o_1 - z_2)^{-d_1^1} & (o_2 - z_2)^{-1} & \cdots & (o_2 - z_2)^{1-d_2^1} & \cdots & (o_{n_1} - z_2)^{-1} & \cdots & (o_{n_1} - z_2)^{1-d_{n_1}^1} \\ \vdots & & \vdots & \vdots & & \vdots & & \vdots & & \vdots \\ (o_1 - z_{n_0})^{-1} & \cdots & (o_1 - z_{n_0})^{-d_1^1} & (o_2 - z_{n_0})^{-1} & \cdots & (o_2 - z_{n_0})^{1-d_2^1} & \cdots & (o_{n_1} - z_{n_0})^{-1} & \cdots & (o_{n_1} - z_{n_0})^{1-d_{n_1}^1} \\ (o_1 - f_2)^{-1} & \cdots & (o_1 - f_2)^{-d_1^1} & (o_2 - f_2)^{-1} & \cdots & (o_2 - f_2)^{1-d_2^1} & \cdots & (o_{n_1} - f_2)^{-1} & \cdots & (o_{n_1} - f_2)^{1-d_{n_1}^1} \\ \vdots & & \vdots & \vdots & & \vdots & & \vdots & & \vdots \\ (o_1 - f_{n_\infty})^{-1} & \cdots & (o_1 - f_{n_\infty})^{-d_1^1} & (o_2 - f_{n_\infty})^{-1} & \cdots & (o_2 - f_{n_\infty})^{1-d_2^1} & \cdots & (o_{n_1} - f_{n_\infty})^{-1} & \cdots & (o_{n_1} - f_{n_\infty})^{1-d_{n_1}^1} \end{vmatrix}$$

which can now be identified as a rational fraction in the z_i , f_i and o_i .

The denominator is easier to compute first: it just consists in the product of the terms $(o_i - z_j)^k$ and $(o_i - f_j)^k$ with the highest power appearing in the determinant, namely $d_i^1 - 1 + \delta_i^1 = \tilde{d}_i^1$ (which is the width of the matrix block corresponding to o_i). Multiplying each line by those factors which appear in it, and the first one by λ , one gets a matrix A with polynomial entries in the z_i , f_i and o_i ; let P be the determinant of that matrix. To compute it, we now forget about the origin of the matrix and note that if any two among $\{z_i\} \cup \{f_i\}$ are equal, then two of the lines of A are equal, hence P is divisible by the product of all the differences between these values.

If on the other hand $o_i = o_j$ for some $i \neq j$, then two blocks of vertical lines become equal, so P is divisible by some power of $(o_i - o_j)$. Getting the value of the power is a little bit more involved but again can be done very explicitly. Start with the following identity (where $a \neq x$ and $k \in \mathbb{Z}_+$):

$$(4) \quad \frac{1}{a-x} = \frac{1}{a} + \frac{x}{a^2} + \cdots + \frac{x^{k-1}}{a^k} + \frac{x^k}{a^k(a-x)}.$$

This can be applied to get an expansion of the first column in the block corresponding to o_j : for the second line, letting $a = (o_i - z_2)$, $x = (o_i - o_j)$ and $k = d_i^1 - 1 + \delta_i^1$ gives an expansion of $(o_j - z_2)^{-1}$ in terms of the entries in the same line in the block of o_i , with coefficients depending only on $(o_i - o_j)$ and hence being the same across all lines, plus a remainder where $(o_i - o_j)^k$ is in factor: more explicitly,

$$(5) \quad \frac{1}{o_j - z_2} = \frac{1}{o_i - z_2} + \frac{o_i - o_j}{(o_i - z_2)^2} + \cdots + \frac{(o_i - o_j)^{k-1}}{(o_i - z_2)^k} + \frac{(o_i - o_j)^k}{(o_i - z_2)^k(o_j - z_2)}.$$

The next columns can be expanded in a similar way, where the expansion of the ℓ -th column in the block of o_j uses the last $(d_i^1 - 1 + \delta_i^1) - (\ell - 1)$ column of the block of o_i and the remainders in the first $(\ell - 1)$ columns of the block of o_j , with coefficients depending only on $(o_i - o_j)$ and a remainder term where the same power of $(o_i - o_j)$ as before factors out. These expansions are obtained from successive derivatives of (4) with respect to x ; the first one being for instance

$$\frac{1}{(a-x)^2} = \frac{1}{a^2} + \frac{2x}{a^3} + \cdots + \frac{(k-1)x^{k-2}}{a^k} + \frac{kx^{k-1}}{a^k(a-x)} + \frac{x^k}{a^k(a-x)^2}.$$

Overall, the power of $(o_i - o_j)$ in P obtained this way is exactly the product of the widths of the corresponding blocks in the matrix.

To summarize, P is divisible by the numerator in the statement of the theorem. Matching the degrees shows that the ratio between $\det J_\Lambda$ and the formula in the statement of the lemma is a constant, which is the main claim. The fact that the determinant is not zero is then a direct consequence of the fact that \mathcal{C}_M is non-degenerate. \square

Proof of Theorem 3. First, we can replace f with e^g in the expression for Φ , and compute the partial derivatives appearing in J_Φ in terms of those appearing in L_Λ . By repeated use of the chain rule, commutativity of derivatives and the fact that \mathcal{C}_M is exact, we get for example

$$\partial_{z_i}[(f_{\lambda,c})^{(j)}(o_k)] = (f_{\lambda,c} \partial_{z_i} g_{\lambda,c})^{(j)}(o_k) = \sum_{\ell=0}^j \binom{j}{\ell} (f_{\lambda,c}^{(\ell)} \partial_{z_i} g_{\lambda,c}^{(j-\ell)})(o_k) = \partial_{z_i}[(g_{\lambda,c})^{(j)}(o_k)]$$

at the point $(\lambda_M, \mathcal{C}_M)$ as soon as $j < d_k^1$. Similar computations in the other variables show that in fact J_Λ and J_Φ are equal at the point $(\lambda_M, \mathcal{C}_M)$ and in particular, $J_\Phi((\lambda_M, \mathcal{C}_M))$ is non-singular.

The remainder of the proof is a completely standard application of the usual Newton algorithm: it suffices to use the smoothness of Φ in all variables to show that $J_\Phi((\lambda, \Phi))$ is non-singular in a neighborhood of $(\lambda_M, \mathcal{C}_M)$, and to expand Ψ at $(\lambda_M, \mathcal{C}_M)$ to obtain a bound of the form

$$\|\Psi((\lambda, \mathcal{C})) - (\lambda_M, \mathcal{C}_M)\| \leq C \|(\lambda, \mathcal{C}) - (\lambda_M, \mathcal{C}_M)\|^2$$

for (λ, \mathcal{C}) close enough to $(\lambda_M, \mathcal{C}_M)$. □

To summarize our construction so far, we are essentially done at the theoretical level: there is a neighborhood U of our point of interest $(\lambda_M, \mathcal{C}_M)$ from which Newton's algorithm is guaranteed to converge, and iterating the subdivision in the first step will bring the approximation within U after finitely many steps. In addition, since the convergence is quadratic, one can iterate a computer implementation until two successive values are indistinguishable within machine precision (or chosen extended precision) at very little cost, meaning that we can get a numerical approximation of $(\lambda_M, \mathcal{C}_M)$ with arbitrarily chosen precision in a reasonable computing time.

A significant issue in practice though is that none of the two steps is quantitative. The speed of convergence of circle packing embeddings to the uniformizing map is not well understood (although our simulations as well as the numerical experiments at the end of [7] suggest that it should be polynomial in the diameter of the smallest circle and exponential in the number of refinements), and the neighborhood U we would be able to explicitly construct by keeping track of all constants implicit in the proof above would certainly be much smaller than the basin of attraction of $(\lambda_M, \mathcal{C}_M)$.

What this means is that the number of subdivisions we should perform in step 1 to be certain to have convergence would make that step computationally unfeasible. We now turn to a way around this issue.

3.3. Step 3: Topological verification. As it turns out, in practice the domain of attraction of Newton's method in the case we are interested in seems to be quite large, at least much larger than continuity arguments for Φ would predict. A natural procedure is therefore to start Newton's algorithm from successive iterations of step 1, starting in fact with no iteration at all, and in each case to see if the iteration converges or not (quadratic convergence when the iteration is successful means that recognizing convergence is very quick).

If the iteration does not converge, or converges to a singular constellation, then we simply refine once more. If it does converge to a non-singular pair $(\lambda_\infty, \mathcal{C}_\infty)$, then we found a constellation with the right signature. It remains to check whether it actually is the one we were looking for (and subdivide once more if it is not).

One can simply do it visually, from a picture of the sign of the imaginary part of the function $f_{\lambda_\infty, \mathcal{C}_\infty}$ like those in the figures of this paper, recovering the triangulation and just checking that it is isomorphic to M . The same thing can probably be automated, although it is not at all clear which algorithms could be proved to work; constructing a region adjacency graph from the picture (thus recovering the dual graph of the triangulation) would be an option, but we did not attempt to implement it.

In practice, except on specially tuned cases designed to test the numerical stability of the whole method (similar for instance to the map shown in Figure 7 with vertices of very large degree), the limit is almost always the right one — though of course this statement is not of a mathematical nature and is merely an empirical observation.

3.4. Step 4: Lattice reduction. At this point, we obtained a numerical approximation of $(\lambda_M, \mathcal{C}_M)$ with the normalization chosen in such a way that the function $f_{\lambda_M, \mathcal{C}_M}$ has a zero at 0, a pole at ∞ and takes the value one at 1. We know in advance that with such normalization, the locations of the other zeros, poles and ones are all algebraic numbers: indeed, there is a choice of normalization such that all locations are algebraic and mapping it to our preferred one can in turn be done by applying a Möbius map with algebraic coefficients.

This means that for each of those algebraic numbers, we are able to obtain an approximation to any precision that we want, in a reasonable computing time. Identifying the minimal polynomial of an algebraic number given such an approximation is a much-studied question, and there exist a variety of classical *integer relation algorithms* to do it, based on lattice-reduction methods. Going into a detailed description of such algorithms is besides the point of this paper; several implementations are freely available, and the results presented below were obtained using one of them (specifically FPLLL, see [2]).

While the empirical observation is that these methods work, again a comment of a more theoretical nature is in order: even though the numbers that we are interested in are all algebraic, the bounds on their degree that one can derive from the proof of Belyi's theorem are enormous and the precision approximation that is required for lattice reduction to provably find the right solution is therefore enormous as well, to the point that making those bounds quantitative is likely to be of little practical use.

A more interesting question is that of the choice of normalization. The one we chose was convenient from the implementation point of view, as it made the formulas in Step 2 explicit, but there is no reason why it would lead to the algebraic numbers of the lowest possible degree, thus compounding the previous remark. Other choices are of course available: for instance one might want the sum with multiplicity of all zeros to be 0 (to get a vanishing coefficient in the numerator of $f_{\lambda_M, \mathcal{C}_M}$). More convincingly, if M has a non-trivial automorphism group, so does its constellation, and one may want (some of) the corresponding automorphisms to be affine transformations of the plane.

This however has little impact on the implementation, as one can always go from one normalization to another after having obtained the approximation in Step 3.

3.5. Step 5: Algebraic verification. The last validation step is of a purely algebraic nature: given the list of locations as roots of integer polynomials, verify that the constellation they form is indeed exact, in which case we have achieved our programme of computing the Belyi function associated to the map M explicitly. This can be rewritten as a collection of algebraic equations that they must satisfy, and can be checked explicitly either by hand (for smaller cases) or by a computer algebra system; again, if the validation failed, it means that the




FIGURE 7. Uniformization of a triangulation of genus 1 with one vertex of high degree (which can be used as a test-case for numerical stability of the algorithm).

polynomials obtained in the previous step are erroneous, and one can re-run the lattice reduction from higher precision approximations, with the guarantee that after finitely many round trips the right solution will be found.

Remark 1. This last step is usually presented as the starting point of the computation of Belyi functions: namely, starting from a hypermap, obtain a system of polynomial equations in the location of the ramification points, and then use elimination theory to solve the system (usually using Gröbner bases). This works well in practice, but does not seem to extend well to higher genus beyond very small maps.

4. NUMERICAL COMPUTATION OF BELYI MAPS: GENUS 1

We now turn to the case of triangulations of the torus, and to elliptic Belyi functions. As mentioned earlier, the main structure of the construction is extremely similar to the one described in the previous section, and we will focus on the differences rather than giving a complete description, using the same notation for corresponding but slightly different objects where it doesn't lead to confusion.

We will work on the space of all non-degenerate constellations of genus 1 and of a given signature defined on a complex torus; however, there is an invariant (the modulus of the torus) which we do not know in advance. This means that the objects we will work with are in fact triples of the form $(\tau, \lambda, \mathcal{C})$ where τ is the modulus of a torus, λ is a normalizing factor and \mathcal{C} is a constellation in the torus $\mathbb{T}_\tau = \mathbb{C}/(\mathbb{Z} + \tau\mathbb{Z})$, and having such a representation of the Riemann surface on which we are working will be very convenient.

So again, fix three tuples of positive integers, $d^0 = (d_i^0)_{1 \leq i \leq n_0}$, $d^\infty = (d_i^\infty)_{1 \leq i \leq n_\infty}$ and $d^1 = (d_i^1)_{1 \leq i \leq n_1}$ with the same sum N , satisfying the genus-1 condition

$$n_0 + n_\infty + n_1 = N,$$

and let $\mathfrak{T} = \mathfrak{T}_{d^0, d^\infty, d^1}$ be the space of all triples $(\tau, \lambda, \mathcal{C})$ where $\tau \in \mathbb{H}$, $\lambda \in \mathbb{C} \setminus \{0\}$ and \mathcal{C} is a non-degenerate constellations on \mathbb{T}_τ with signature (d^0, d^∞, d^1) . \mathfrak{T} is a complex manifold of dimension $N + 2$. We will always see \mathbb{T} as the quotient of the complex plane by the lattice $\mathbb{Z} + \tau\mathbb{Z}$, represent the locations of the stars of \mathcal{C} as complex numbers and identify \mathfrak{T} as an open subset of \mathbb{C}^{N+2} ; if $\mathcal{C} = (Z, P, O) \in \mathfrak{C}$, let

$$f_{\tau, \lambda, \mathcal{C}} : z \mapsto \lambda \frac{\prod \zeta(z - z_i)^{d_i^0}}{\prod \zeta(z - p_i)^{d_i^\infty}}$$

be defined as above, and consider the map $\Phi : \mathfrak{T} \rightarrow \mathbb{C}^N$ defined by

$$\Phi((\tau, \lambda, \mathcal{C})) := (f(o_1) - 1, f'(o_1), \dots, f^{(d_1^1-1)}(o_1), f(o_2) - 1, f'(o_2), \dots, f^{(d_{n_1}^1-1)}(o_{n_1})),$$

where to lighten notation we let $f = f_{\tau, \lambda, \mathcal{C}}$.

To normalize the map embeddings and ensure that f is doubly periodic with periods 1 and τ (or in other words, to have it defined on the torus \mathbb{T}_τ), let \mathfrak{T}^* be the set of all triples $(\tau, \lambda, \mathcal{C}) \in \mathfrak{T}$ such that

$$(6) \quad \sum d_i^0 z_i = \sum d_i^\infty p_i = 0.$$

\mathfrak{T}^* is a manifold of complex dimension N which can be seen as an open subset of \mathbb{C}^N by listing the locations of all the stars except for z_1 and p_1 ; the restriction $\Phi^* : \mathfrak{T}^* \rightarrow \mathbb{C}^N$ can thus be seen as a map from an open subset of \mathbb{C}^N to \mathbb{C}^N , which is clearly analytic in all its variables. We are interested in finding one specific preimage of $(0, \dots, 0)$ by Φ , out of finitely many.

We are now exactly in the same situation as before, the only difference being the definition of f , so we will briefly review the relevant changes needed to the previous steps.

4.1. Step 1: Approaching the complex structure. Here the construction starts from a triangulation T of genus 1, and again there exists a unique $\tau_T \in \mathbb{H}$ and an essentially unique circle packing on the torus of modulus τ_T having the combinatorics of T . Equivalently, the universal cover \hat{T} of T is a doubly periodic triangulation of the plane, and there is an essentially unique locally finite circle packing of the plane with the combinatorics of \hat{T} ; this circle packing is automatically doubly periodic, and can be normalized in such a way that it has the two periods $\{1, \tau_T\}$.

We want to define the starting point for the Newton iterative scheme using τ_T as the modulus and the locations of the centers of the circles as those of the stars; and here as well, under successive refinements, these converge to the constellation corresponding to M . One slight problem is that there is no reason why the normalization condition (6) can be made

to hold by a suitable choice of embedding of a circle packing — but we can at least ensure that $\sum d_i^0 z_i = 0$, which makes the embedding unique, and simply ignore the sum of the $d_i^\infty p_i$ (which will tend to 0 as the triangulation is refined more and more).

4.2. Step 2: Newton’s method. This part works exactly the same way as before, with

$$\Psi((\tau, \lambda, \mathcal{C})) := (\tau, \lambda, \mathcal{C}) - J_\Phi((\tau, \lambda, \mathcal{C}))^{-1} \Phi((\tau, \lambda, \mathcal{C})),$$

and the only point to check is the non-singularity of the Jacobian matrix at the point $(\tau_M, \lambda_M, \mathcal{C}_M)$. This can be done in a very similar way as on the sphere, as far as the variables besides τ are concerned: for fixed τ the Jacobian determinant is an elliptic function in all its other variables, and it factorizes for the same reason as in the rational case. Alternatively, one can also get the non-singularity of the Jacobian by moving all the locations of the points closer and closer to the origin (while satisfying the constraints listed above), and noticing that the asymptotic behavior of the determinant, once properly normalized, is given by the case of genus 0.

The variable τ is a bit more problematic, because the derivatives in τ of the functions involved in Φ are not as explicit as the derivatives in z . Instead, one way to go around the issue is to argue that a different way of proceeding, from the data of a toroidal triangulation, is as follows: take a large square of $N \times N$ periods in the universal cover of M , and stitch it with a copy of itself along the boundary of the square to obtain a triangulation of the sphere. This triangulation can be uniformized (for instance using the method in the previous section) and the local behavior of the uniformizing function near the center of the square, as $N \rightarrow \infty$, becomes periodic and asymptotic to the uniformizing map of M itself. In particular, the fact that the Jacobian is non-degenerate in genus 1 can be extracted from the corresponding statement in genus 0.

Remark 2. While that last remark gives an alternative strategy to solve our initial problem as well, it has two main drawbacks: first, it is difficult to estimate the level of precision that one would be obtained as a function of N (presumably it would be polynomial in N); second, and very related, is that to get the kind of approximation needed to apply the next step, the value of N would have to be taken so large that there would be no hope of actually implementing the programme. It would still be a possibility to obtain the starting point of step 2 in this way, thus replacing step 1; but we did not try this route.

4.3. Steps 3, 4 and 5: Identification and validation. Here, not much needs to be changed at all except for notation; the questions raised, whether the successive objects that we construct correspond to the one we are looking for, are the same, and the method is exactly parallel to that used in genus 0.

One point needs to be made though. In the sphere, the freedom in the choice of embedding meant that finding the constellation leading to algebraic numbers with minimal degree was an issue. Here, there is much more rigidity from the conditions (6), which itself is very natural, and there is no choice at all, beyond the usual discrete $SL_2(\mathbb{Z})$ action, in the modulus τ ; indeed the elliptic curve defined by M is uniquely defined. This is quite convenient, especially since determining τ was our primary goal in the case of genus 1.

5. A FEW EXAMPLES

The front page of this paper shows the Belyi function of the tripartite refinement of a uniformly sampled random triangulation of the sphere with 15 vertices. We validated the


algorithm of a few known cases, such as those listed in [1], but chose to focus on genus 1 for the examples of this section.

We list here the Belyi functions obtained from all the triangulations of the torus with up to 3 vertices and all degrees at least equal to 3. In each case, the triangulation is refined into its tripartite refinement and we list the hypermap description of the original triangulation, a graphical representation of its combinatorics, and then the sign of the imaginary part of the Belyi function and in a few instances the exact value of the associated j -invariant (or its minimal polynomial in case that is more useful). For triangulations with 4 vertices we give one example in the same form, and only the graph of the covering map for the others.

In several cases the original triangulation is itself tripartite, so the representation is not minimal; but the value of j is not affected by the refinement, and the dessin associated to the initial triangulation is a subset of that of the refined one.



Each finite triangulation T has countably many coverings $(T_{k,l})$ that are themselves finite triangulations, having as fundamental domains unions of finitely many copies of the fundamental domain of T arranged as a $k \times l$ rectangle. Those are omitted in the catalog below, and only the minimal one is listed.

5.1. Size 1. There is only one triangulation of the torus with one vertex, and its natural embedding is the usual triangular lattice composed of equilateral triangles:





5.2. Size 2. In addition to the double coverings of the previous one, there are two minimal triangulations of the torus with two vertices: one is a refinement of the triangular lattice (with the same j invariant equal to 0) and the other is the face-centered square lattice.

$$\begin{aligned}\sigma &= (0785219104)(3116) \\ \alpha &= (01)(23)(45)(67)(89)(1011) \\ \varphi &= (026)(148)(3510)(7119)\end{aligned}$$





$$\begin{aligned}\sigma &= (07104)(196311852) \\ \alpha &= (01)(23)(45)(67)(89)(1011) \\ \varphi &= (026)(148)(3510)(7911)\end{aligned}$$





5.3. Size 3. There are 9 triangulations of the torus with 3 vertices and minimal degree at least 3 that are not coverings of previously displayed cases. They turn out to all have rational j -invariants.

$$\begin{aligned}\sigma &= (07171495126315104)(182)(111613) \\ \alpha &= (01)(23)(45)(67)(89)(1011)(1213)(1415)(1617) \\ \varphi &= (026)(149)(3814)(51013)(71216)(111517)\end{aligned}$$




$$\begin{aligned}\sigma &= (0716104)(1912631114852)(131517) \\ \alpha &= (01)(23)(45)(67)(89)(1011)(1213)(1415)(1617) \\ \varphi &= (026)(148)(3510)(71217)(91413)(111615)\end{aligned}$$




$$\sigma = (07\ 14\ 10\ 4)(19\ 17\ 15\ 13\ 8\ 5\ 2)(3\ 11\ 16\ 12\ 6)$$

$$\alpha = (01)(23)(45)(67)(89)(10\ 11)(12\ 13)(14\ 15)(16\ 17)$$

$$\varphi = (02\ 6)(14\ 8)(35\ 10)(7\ 12\ 15)(9\ 13\ 16)(11\ 14\ 17)$$




$$j = -3072$$




$$\sigma = (05\ 13\ 10\ 6\ 2\ 17\ 14\ 8\ 4\ 3)(9\ 16\ 12)(11\ 17\ 15)$$

$$\alpha = (01)(23)(45)(67)(89)(10\ 11)(12\ 13)(14\ 15)(16\ 17)$$

$$\varphi = (02\ 4)(13\ 6)(58\ 12)(7\ 10\ 15)(9\ 14\ 17)(11\ 13\ 16)$$




$$j = -3072$$




$$\sigma = (07\,17\,15\,10\,4)(19\,16\,13\,14\,8\,5\,2)(3\,11\,12\,6)$$

$$\alpha = (01)(23)(45)(67)(89)(10\,11)(12\,13)(14\,15)(16\,17)$$

$$\varphi = (02\,6)(14\,8)(35\,10)(7\,12\,16)(9\,14\,17)(11\,15\,13)$$




$$j = \frac{35\,152}{9}$$




$$\sigma = (05\,13\,14\,8\,4\,3)(17\,17\,15\,10\,6\,2)(9\,16\,11\,12)$$

$$\alpha = (01)(23)(45)(67)(89)(10\,11)(12\,13)(14\,15)(16\,17)$$

$$\varphi = (02\,4)(13\,6)(58\,12)(7\,10\,16)(9\,14\,17)(11\,15\,13)$$




$$j = -\frac{33\,268\,701}{256}$$




$$\sigma = (07\ 16\ 8\ 2\ 1\ 11\ 14\ 9\ 4)\ (3\ 15\ 12\ 6)\ (5\ 17\ 13\ 10)$$

$$\alpha = (01)\ (23)\ (45)\ (67)\ (89)\ (10\ 11)\ (12\ 13)\ (14\ 15)\ (16\ 17)$$

$$\varphi = (02\ 6)\ (14\ 10)\ (38\ 14)\ (59\ 16)\ (7\ 12\ 17)\ (11\ 13\ 15)$$




$$j = \frac{8\ 429\ 568}{15\ 625}$$




$$\sigma = (07\ 14\ 8\ 5\ 2\ 19\ 16\ 10\ 4)\ (3\ 11\ 12\ 6)\ (13\ 17\ 15)$$

$$\alpha = (01)\ (23)\ (45)\ (67)\ (89)\ (10\ 11)\ (12\ 13)\ (14\ 15)\ (16\ 17)$$



$$\varphi = (02\ 6)\ (14\ 8)\ (35\ 10)\ (7\ 12\ 15)\ (9\ 14\ 17)\ (11\ 16\ 13)$$



$$j = -\frac{1\ 636\ 015\ 539}{41\ 229\ 056}$$





$$\begin{aligned}\sigma &= (05\ 13\ 14\ 11\ 8\ 4\ 3)(17\ 15\ 16\ 10\ 6\ 2)(9\ 17\ 12) \\ \alpha &= (01)(23)(45)(67)(89)(10\ 11)(12\ 13)(14\ 15)(16\ 17) \\ \varphi &= (02\ 4)(13\ 6)(58\ 12)(710\ 14)(9\ 11\ 16)(13\ 17\ 15)\end{aligned}$$



5.4. Size 4. With 4 vertices there are already too many triangulations to make listing them very useful. We still display one in detail, and will only show the picture for the others:


$$\begin{aligned}\sigma &= (07\ 19\ 22\ 14\ 9\ 4)(111\ 21\ 23\ 16\ 8\ 2)(315\ 20\ 12\ 6)(517\ 18\ 13\ 10) \\ \alpha &= (01)(23)(45)(67)(89)(10\ 11)(12\ 13)(14\ 15)(16\ 17)(18\ 19)(20\ 21)(22\ 23) \\ \varphi &= (02\ 6)(14\ 10)(38\ 14)(59\ 16)(712\ 18)(11\ 13\ 20)(15\ 22\ 21)(17\ 23\ 19)\end{aligned}$$





This is the triangulation used as an example in [4, 12], where the elliptic curve is derived formally. It is defined over $\mathbb{Q}[\sqrt{7}]$ and one can check that the j -invariant obtained by our method is the right one (the discriminant of the polynomial above is $7 \cdot (345\,128)^2$). Note though that the proof as detailed in [4] is 4 pages long and moreover relies very strongly on the additional symmetries of the triangulation — in the picture above one can readily see that the embedding is symmetric under reflection by the line going through 0 and $1 + \tau$ and to another one orthogonal to it, which both correspond to automorphisms of order 2 of the triangulation.


This is a strong indication that very small examples are computationally difficult to address, and slightly larger ones, or even cases of small size but no symmetry, are beyond these methods.

In comparison, the programme described here gets the exact constellation to machine precision (*i.e.*, to within 10^{-13}) in less than a tenth of a second and producing enough digits to obtain the minimal polynomial for j takes of the order of 20 seconds on a standard laptop.









REFERENCES

- [1] N. M. ADRIANOV, N. Y. AMBURG, V. A. DREMOV, Y. Y. KOCHETKOV, E. M. KREINES, Y. A. LEVITSKAYA, V. F. NASRETDINOVA, AND G. B. SHABAT, *Catalog of dessins d'enfants with no more than 4 edges*, J Math Sci, 158 (2009), pp. 22–80.
- [2] M. ALBRECHT, D. CADÉ, X. PUJOL, AND D. STEHLÉ, *fpLL-4.0, a floating-point LLL implementation*. Available at <http://perso.ens-lyon.fr/damien.stehle>.
- [3] L. BARTHOLDI, X. BUFF, H.-C. GRAF VON BOTHMER, AND J. KRÖKER, *Algorithmic construction of hurwitz maps*, Exp. Math., 24 (2015), pp. 76–92.
- [4] M. BAUER AND C. ITZYKSON, *Triangulations*, Discr. Math., 156 (1996), pp. 29–81.
- [5] J. BÉTRÉMA AND A. ZVONKIN, *La vraie forme d'un arbre*, in TAPSOFT '93: theory and practice of software development (Orsay, 1993), vol. 668 of Lecture Notes in Comput. Sci., Springer, Berlin, 1993, pp. 599–612.
- [6] F. BEUKERS AND H. MONTANUS, *Explicit calculation of elliptic fibrations of K3-surfaces and their Belyi-maps*, in Number theory and polynomials, J. McKee and C. Smyth, eds., Cambridge Univ. Press, Cambridge, Cambridge, 2008, pp. 33–51.
- [7] P. L. BOWERS AND K. STEPHENSON, *Uniformizing dessins and Bely maps via circle packing*, Mem. Amer. Math. Soc., 170 (2004), pp. xii–97.
- [8] J. W. CANNON, *The combinatorial Riemann mapping theorem*, Acta Math., 173 (1994), pp. 155–234.
- [9] E. GIRONDO AND G. GONZÁLEZ-DIEZ, *Introduction to compact Riemann surfaces and dessins d'enfants*, vol. 79 of London Mathematical Society Student Texts, Cambridge University Press, Cambridge, 2012.

- [10] F. KLEIN, *Vorlesungen über das Ikosaeder und die Auflösung der Gleichungen vom fünften Grade*, Birkhäuser Verlag, Basel; B. G. Teubner, Stuttgart, 1993. Reprint of the 1884 original, Edited, with an introduction and commentary by Peter Slodowy.
- [11] S. K. LANDO AND A. K. ZVONKIN, *Graphs on surfaces and their applications*, vol. 141 of Encyclopaedia of Mathematical Sciences, Springer-Verlag, Berlin, 2004. With an appendix by Don B. Zagier, Low-Dimensional Topology, II.
- [12] G. B. SHABAT AND V. A. VOEVODSKY, *Drawing curves over number fields*, Progr. Math., 88 (1990), pp. 199–227.
- [13] D. SINGERMAN AND R. I. SYDDALL, *Belyi uniformization of elliptic curves*, Bull. London Math. Soc., 29 (1997), pp. 443–451.
- [14] K. STEPHENSON, *Introduction to Circle Packing*, Cambridge University Press, 2005.
- [15] A. ZVONKIN, *Belyi functions: examples, properties, and applications*, in Proceedings AAECC, 2008.

UMPA, ENS LYON AND HCM, BONN

E-mail address: vbeffara@ens-lyon.fr

URL: <http://perso.ens-lyon.fr/vincent.beffara/>